

# Demystifying AI for Business

Artificial intelligence (AI) and machine learning (ML) are helping businesses around the world improve decision-making, enhance customer engagement, and increase their bottom lines. A thorough understanding of the concepts and key terms is critical as you consider what AI and ML can do for your business.

## Introduction: AI all around us

Businesses today have more data than ever before, and that data houses invaluable insights. This treasure trove of data has the potential to accelerate decision-making, improve forecasting, reduce operational overhead, predict equipment failures before they happen, and overall, gain a competitive advantage in an increasingly global marketplace. But how? AI and ML give you the tools you need to harness your businesses data and turn it into powerful insights. And whether you realize it or not, you have been using these tools almost every day.

Artificial intelligence is enabled by machine learning. Machine learning represents the tools, techniques, and technologies that enable developers to build products or automate processes that include human-like capabilities. A machine learning system can accurately predict or understand something based on data it is given.

The applicability of this technology is broad. Contact call centers use automated chatbots to improve customer service and reduce wait-times, manufacturers use pattern analysis across billions of inventory data points to provide accurate supply chain forecasts, broadcasters use real-time statistics and play predictions in your favorite sports, and the images you see in your social media feed and the recommendations when you shop online are all driven by AI and ML.

AI and ML have been around for decades, but

they are experiencing a strong surge of interest recently among business thought leaders. This is because the amount of data for analysis is growing and the tools to store, analyze, and build ML systems are becoming more accessible. The cloud, which provides purpose-built infrastructure — based on leading-edge hardware — plus the raw computing power necessary to run these complex analyses, is the key.

In 2018, Gartner predicted that the AI market will be worth [\\$1.2 trillion](#), indicating that that around half of companies surveyed are [planning to adopt](#) some form of AI and ML solution in the near future.

While it is clear that AI and ML are important tools for businesses looking forward, there is a big gap in the skills, training, and overall know-how that business leaders are facing today. A recent [Spiceworks survey](#) shows that 50 percent of businesses have not implemented AI due to a lack of use cases in their organization. To get started, let's take a look at the key concepts and how to start the process of applying AI and ML in your company.

## How does AI work?

Artificial intelligence is powered by machine learning technology. So, to understand AI you need to click down into ML. A key aspect of ML involves 'teaching' compute systems how to make predictions (also called inferences) based on given parameters.

Starting with an algorithm, or set of rules, the system is 'trained' using massive amounts of

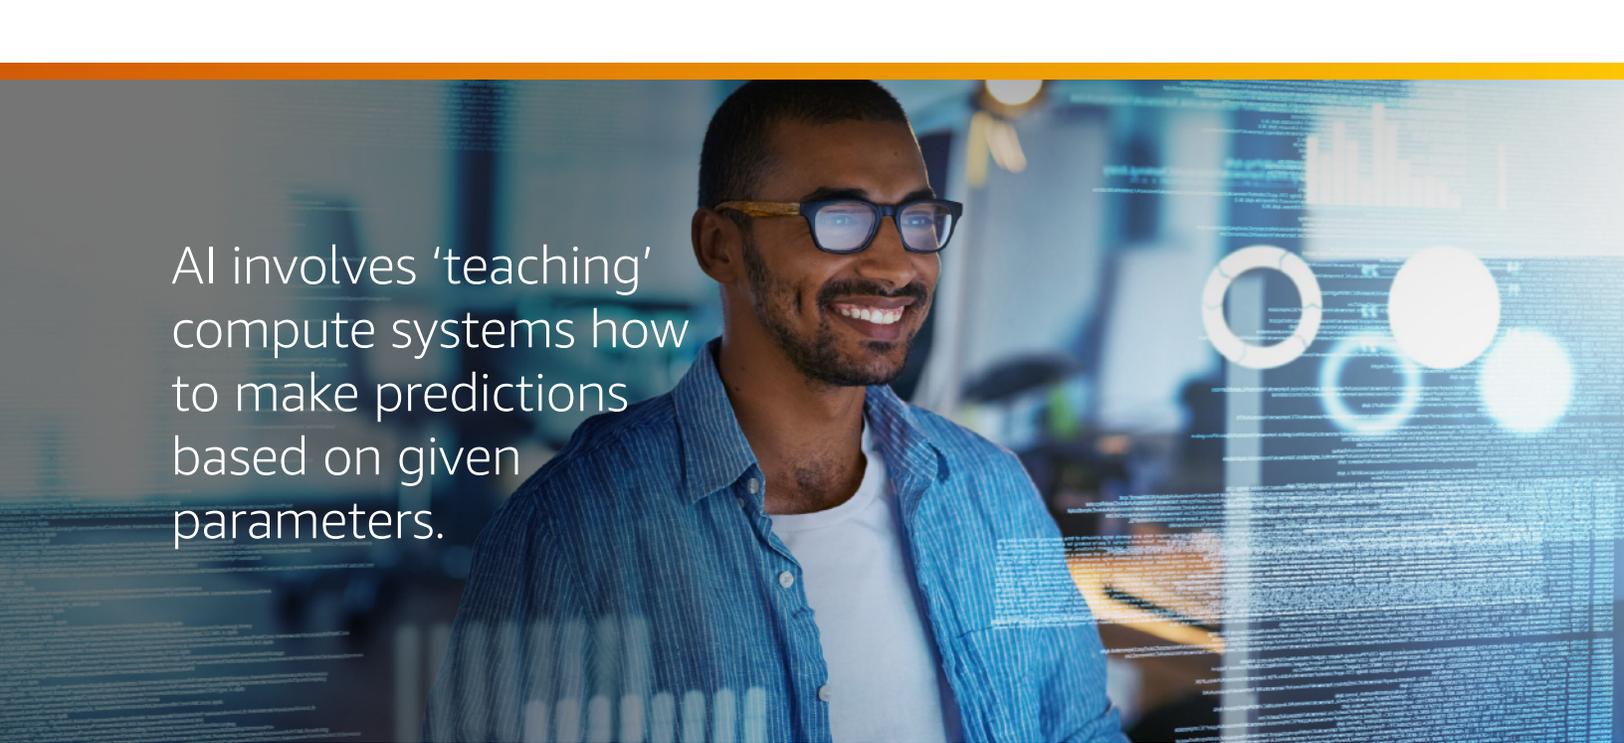
## Common uses of AI and ML include:

**Retail operations:** AI helps retailers predict sales performance based on historical trends and real-time tracking of product churn. It can automatically alert store managers to restock shelves, without any need for human inventory management.

**Online advertising:** Marketers use AI to learn and classify audience preferences, so they can predict which kinds of ads will get the most clicks and refine that model over time.

**Healthcare:** Medical professionals use AI to find trends in electronic health records, analyzing risk factors like the frequency of ER visits correlated by age and weight, for example — and then identifying patients at risk and improving their outcomes.

**Financial Services:** AI learns from millions of transaction records how to detect specific kinds of fraud, and banks are using it to determine loan eligibility for new applicants based on payment and default information gathered from years of application data.



AI involves 'teaching' compute systems how to make predictions based on given parameters.

existing data. This machine learning is similar to the process of human learning. Think about this in the context of how students read a textbook. They consume the data in a textbook, like mathematical rules, for example, to understand how to solve math problems. Then they apply those principles and learning to new information. Doctors are another example. They use the information they absorbed in medical school to design and predict the efficacy of treatment protocols, even for conditions they didn't study specifically.

To bring this to life, let's take a look at sentiment analysis. Let's say a company wants to measure customer satisfaction with a new product or service, and many of the customers are talking about the product on social networks. An algorithm may be written to detect the presence of certain words (great, awesome, terrible, miserable), symbols (smiley or angry faces), and punctuation marks. Some words are clearly positive, others negative, and others may be neutral, depending on context. By feeding the system thousands upon thousands of social media posts and classifying the data accordingly (positive, negative, neutral), it will 'learn' what patterns to look for in future, unclassified posts and then predict the sentiment of the post automatically.

Once trained, the function that accepts new data as an input and returns a classification as output is called a model. With enough training data and compute power, ML models can learn complex language quirks like double meanings and sarcasm, ultimately achieving high levels of accuracy. Sentiment analysis has many uses, from customer experience management to market research.

Another emerging set of use cases involves the Internet of Things (IoT) — the growing number of sensors, cameras, and other devices that are collecting data and transmitting it to repositories. This data could take the form of diagnostic readings from industrial equipment, video feeds from surveillance cameras, and spoken requests to voice-recognition systems.

Every day, IoT devices generate massive amounts of data. ML is used to analyze the data and predict equipment failure and maintenance needs, anticipate customer requests, and measure vehicle traffic. Organizations use these predictions to aid decision-making, automate processes, avoid costly problems, and gain a [competitive advantage](#) in their fields.

These are just two good examples of AI and ML in practice, there are many approaches that can be taken to gain insights using AI and ML across every business function or industry vertical.

### Amazon Web Services eases your entry to AI and ML

AI and ML adoption is still in its early stages, and companies are working through the challenges of understanding how these technologies can benefit their business, finding the right analytics solutions, setting up the right infrastructure, and sourcing the right data science expertise to implement the technology.

As we mentioned, cloud computing offers a compelling alternative to on-premises infrastructure when it comes to machine learning, because of its near-infinite scalability, flexibility, and readily available computational power.

AI IN ACTION

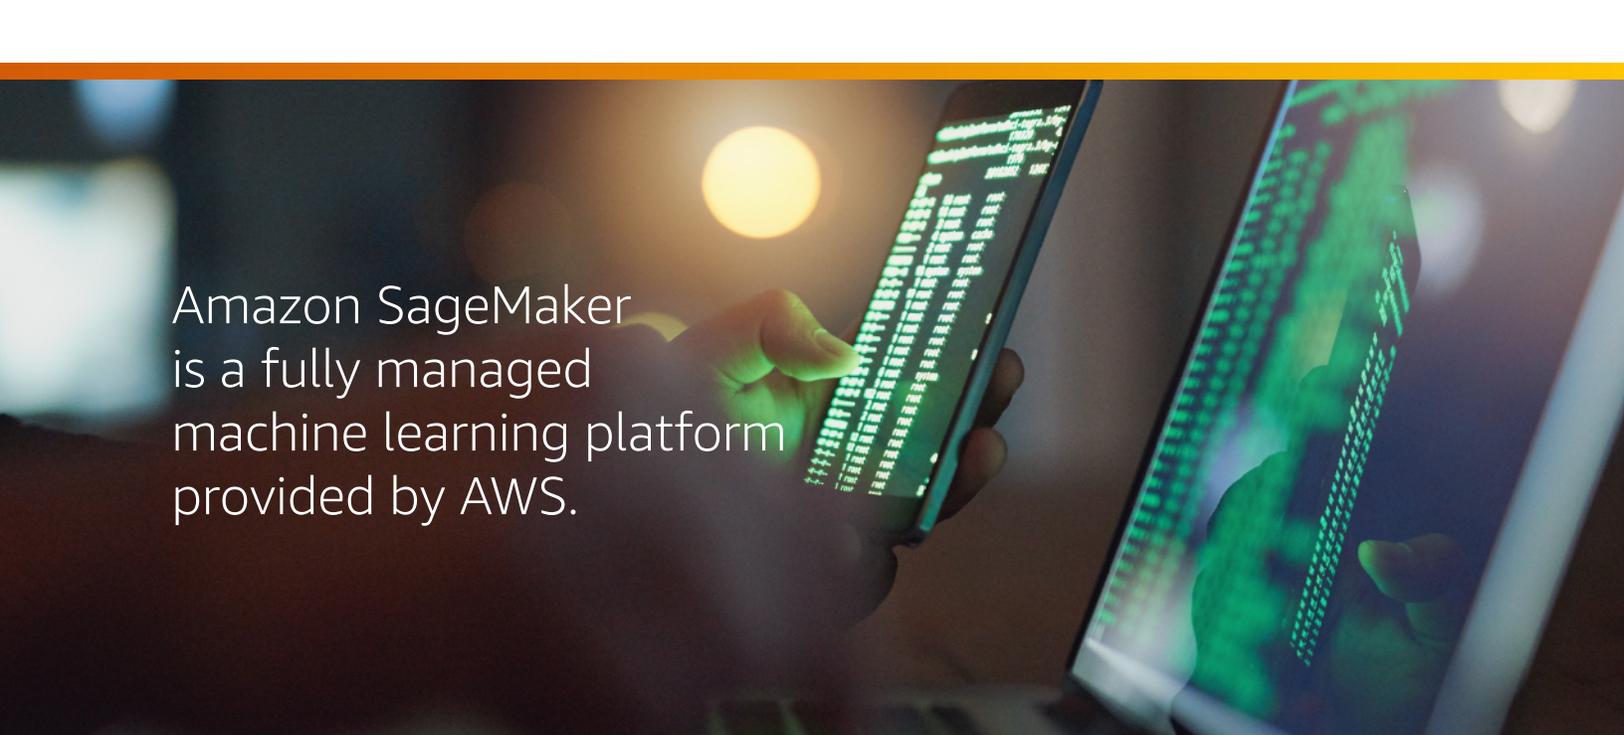


### Electronic Caregiver detects activity using AI

Electronic Caregiver takes the traditional 'medical alert' pendant to the next level, with a wearable wristband that doubles as an emergency alert button, and a location-enabled pocket device, so responders know exactly where a subscriber is and have access to their medical history.

The solution uses AWS services such as Amazon SageMaker to track activity and ensure that subscribers are following healthy routines. Any prolonged period of inactivity will trigger a potentially life-saving alert. AWS language-recognition tools round out the Electronic Caregiver package, connecting subscribers with advanced response teams.

More machine learning is built on Amazon Web Services (AWS) than anywhere else. And by sharing our AI expertise and machine learning capabilities, AWS provides a smooth and simple on-ramp for companies looking to embrace AI and ML. AWS supports a wide range of ML



# Amazon SageMaker is a fully managed machine learning platform provided by AWS.

frameworks and tools in the cloud, and at the edge on your connected IoT devices. These include pre-built algorithms and models, with optimized infrastructure — running on custom Intel Xeon Scalable processors — designed for performance when running compute-intensive machine learning workloads.

Let's take a closer look at some of the machine learning services that AWS offers to power AI and ML for your business:

**Amazon SageMaker:** Machine learning has been challenging for most developers, because the process to build and train models, and then deploy them into production, is inherently complicated and slow. Amazon SageMaker is a fully managed machine learning platform provided by AWS, which removes the complexity that holds back developer success. Amazon SageMaker includes modules that can be used together or independently to build, train, and deploy your machine learning system. It includes algorithm and model authoring tools, simplified connections to internal and cloud-based data sources, and a library of pre-built algorithms that are optimized to run on cloud resources. Simply put, SageMaker is the tool that will allow your existing developers to become machine learning developers.

**AWS Greengrass:** IoT is transforming the world we live in. Across industries, in both the public and private sector, IoT is connecting people and making data more accessible. AWS Greengrass is software that lets you run local compute, messaging, data caching, sync, and ML inference capabilities for connected devices in a secure way. Meaning, you can not

only connect and manage your devices easily, you can also generate meaningful insights. Retailers, cruise lines, and amusement parks are investing in IoT applications to provide better customer service. For example, you can run object detection models at amusement parks to keep track of visitor count. Cameras locate the visitors and maintain a running headcount locally without having to send massive amounts of video feed to the cloud. This solution can predict wait times at popular theme park rides and help improve the customer experience. AWS Greengrass, running on Intel technology, delivers a secure, intelligent 'edge' that allows developers to create new applications easily from edge to cloud.

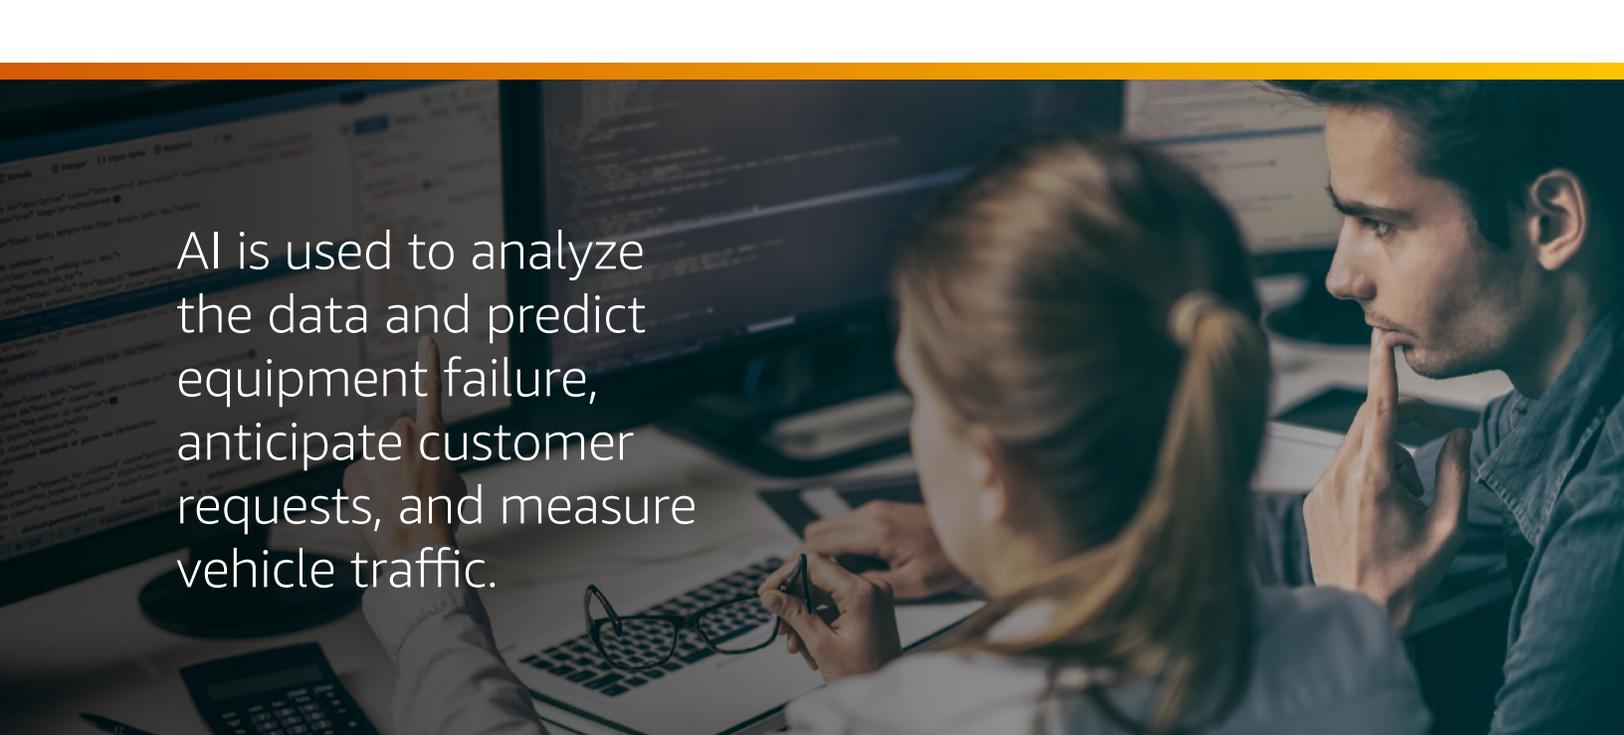
**Amazon Elastic Compute Cloud (EC2):** Defining how AI and ML can benefit your business is just one part of the equation. Building and running the computing infrastructure with the power needed to support ML applications is complex and costly. Amazon EC2 is a web service that takes these pains away. EC2 provides secure, resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers, and it offers a variety of compute instances optimized for many ML use cases. Amazon EC2 C5 instances, for example, run on highly customized and powerful Intel Xeon Scalable Processors and bring great price for performance value for training a variety of ML models and inference functions. AWS also offers GPU-based instances for ML training.

AI IN ACTION  **grammarly**

## AI edits writing at Grammarly

Grammarly is an AI-powered digital writing assistant used by millions of people every day to make their communication clear and effective. Along with spelling and grammar, Grammarly provides feedback on writing issues such as wordiness, vagueness and hedging, word choice, sentence structure, and plagiarism. The service 'reads' and understands written text automatically using machine learning. When you write with Grammarly, its AI analyzes each sentence and looks for ways to improve it, whether it's correcting a verb tense, suggesting a stronger synonym, or offering a clearer sentence structure.

Part of Grammarly's product suite includes the Grammarly Keyboard, a mobile product. On mobile, Grammarly accomplishes some of its tasks with models built and trained using AWS SageMaker. Grammarly runs inferences on SageMaker in the cloud using the Amazon EC2 C5 instance, and it deploys some models that run on mobile devices even when they're not connected to the Internet — a prime example of edge computing.



AI is used to analyze the data and predict equipment failure, anticipate customer requests, and measure vehicle traffic.

## Getting started with AI and ML

You've already taken a step forward in getting started by reading this tech brief to get familiar with AI and ML, but what's next? Here are three actions you can take to move your business forward on its AI and ML journey:

### 1. Assess your business capabilities

Whether your business goals are improving customer service, reducing waste in the supply chain, or accelerating your innovation pipeline, think about where you are missing insights that could help you achieve these goals and transform your business. Having the vision of what you want your business to achieve through AI and ML is just the start. What you want to accomplish and what you have the operational and organizational ability to actually achieve are likely quite different. You'll need robust data sets and at least some foundational ML knowledge inside the organization.

### 2. Build the skills needed inside your organization

If you want to expand the machine learning knowledge and skills within your organization, AWS DeepLens is the easiest way to get started. The world's first deep learning-enabled video camera enables developers to learn the basics of machine learning by exploring practical computer vision use cases,

including pre-built models for face detection, object detection, activity detection, and more. DeepLens is able to run inference models on the device at the edge using Intel Atom processors. Your developers can get started with machine learning in as little as 10 minutes using the tutorials, code, and pre-trained models that are packaged with the device.

### 3. Ask for help

Once you are clear on your goals and capabilities, you're ready to start building and iterating. You can start with a small pilot project and go from there, or go all-in on a big idea. Either way, help is at hand from AWS. You can talk to our experts at the Amazon Machine Learning Solutions Lab, a program that will help you accelerate your use of machine learning in products and processes. The ML Solutions Lab provides you access to the talent that built many of Amazon's machine learning-powered products and services to accelerate your own machine learning expertise through direct guidance and instruction.

To learn more about AWS powered by Intel, please visit [ml.aws](https://ml.aws)

## GLOSSARY OF AI TERMS

**Artificial intelligence:** Systems or computers that can accurately predict or understand something based on data they're given.

**Machine learning:** A sub-set of AI. The process of developing systems that enable AI -- the techniques and technology that give computers the ability to learn.

**Algorithm:** A set of detailed instructions for solving a problem or completing a task. Algorithms enable computers to make suggestions using data.

**Model:** A trained algorithm that can analyze new data and make predictions based on what it has learned.

**Training:** The process of teaching the compute system how to perform the actions you want. Machine learning models are trained with large sets of historical data in order to improve accuracy.

**Inference:** The output of a trained model. Also called predictions or recommendations.

**IoT:** Internet of Things. A network of connected devices that collect and exchange data.

**Edge computing:** The process of performing machine learning directly on IoT devices that are collecting data.